WHITE PAPER

# SUPERMICRO DLC-2 ARCHITECTURE REDUCES DATA CENTER POWER, SPACE, WATER, AND COSTS

*A Range of Innovative Solutions Reduces the PUE, WUE, and SUE of Entire Data Centers Through a Combination of New Liquid-Cooled Optimized Servers and Cooling Infrastructure Improvements*

TABLE OF CONTENTS


Introduction . . . . . 1
Summary of Supermicro DLC-2 . . . . . 2
Why Liquid Cooling? . . . . . 2
4 Measures of Liquid Cooling Effectiveness . . . . . 3
DLC-2 Data Center Improvements . . . . . 3
Conclusion . . . . . 5
For More Information . . . . . 6


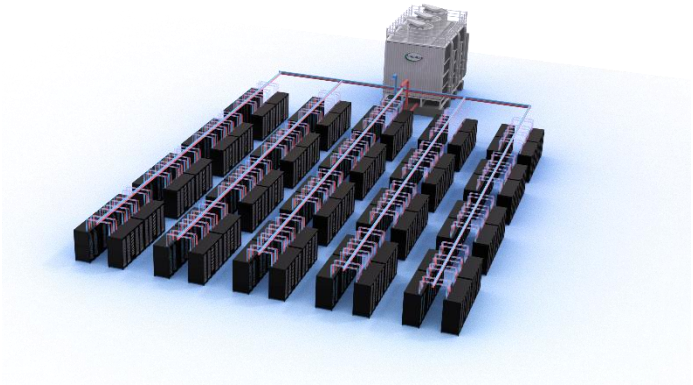

## Introduction

As data centers strive to meet rising computing power and efficiency demands, liquid cooling has become a critical infrastructure necessity for the entire facility. Liquid cooling solutions offer optimal cooling efficiency to support the new generation of GPUs and CPUs compared to traditional air-cooled options. The Supermicro Data Center Building Block Solutions (DCBBS) enhances modern liquid-cooled data centers by saving power, reducing footprint, and decreasing water consumption, all leading to lower operating costs for our clients. To quantify this, four key measurements are critical to consider when designing and deploying a modern data center or retrofitting an existing one: PUE (Power Usage Effectiveness), SUE (Space Usage Effectiveness), WUE (Water Usage Effectiveness), and $UE (Money Usage Effectiveness or MUE). Each measurement is essential for assessing overall cost and power savings when building new data centers or upgrading existing ones.

## Summary of Supermicro DLC-2 (Direct-to-Chip Liquid Cooling)

While liquid-cooling is becoming an accepted method to cool the latest rack-scale solutions, the heat generated by the combination of GPUs, CPUs, and other components is accelerating faster than existing solutions can address. With each new generation of both GPUs and CPUs, the TDPs are rising, and although more performance/watt is growing, there is a clear need to create even more heat removal from the server, rack, and data center. Supermicro's DLC-2 consists of several innovative technologies & benefits, which build upon and enhance the first-gen DLC solutions.

- Supporting a higher inlet temperature for the server coolant
- Fan Speed & air flow reduction
- Significantly reduced noise levels
- Additional cold plate coverage of components in the server
- New hybrid cooling towers
- Increased Capacity of in-rack Coolant Distribution Units (CDUs)

## Why Liquid Cooling?

The increase in the TDP for both GPUs and CPUs has drastically changed the data center demand for more power at the server, rack, and facility levels. Liquid cooling is much more effective at removing heat from hot surfaces, particularly GPUs and CPUs, when compared to airflow over the chip. In liquid cooling, the supply liquid (ambient temperature water) flows over the hot component, transferring heat to the liquid. An external water tower system must cool the resulting warmer liquid (which can be ambient temperature plus 5 to 10 degrees Celsius). The cooled liquid is then returned to the servers, and the cycle continues.

### Supermicro Liquid Cooling Advantages

Reduce Power Consumption
Reduce Data Center Size
Reduce Water Usage

Reduce Overall Cost

## 4 Measures of Liquid Cooling Effectiveness

The effectiveness of liquid cooling can be measured in several ways. A recognized measurement is the PUE (Power Usage Effectiveness). While the standard definition is the power ratio in the data center compared to that used by IT equipment, a more useful metric would be to compare the power used in a liquid-cooled data center to an air-cooled data center.

Assuming the same performance from both liquid and air-cooled data centers, we are calculating the effectiveness of data centers using our DLC-2 GPU servers as follows:

$$PUE_{DLC\ vs\ AC} = \frac{Power\ for\ Liquid\ Cooled\ Solution}{Power\ for\ Air\ Cooled\ Solution} = \frac{IT\ Power[4U\ LC\ DLC\ System] \times 1.1_{DLC\ PUE}}{IT\ Power[10U\ AC\ System] \times 1.5_{AC\ PUE}} = 60\%$$

$$SUE_{DLC\ vs\ AC} = \frac{Space\ for\ Liquid\ Cooled\ Solution}{Space\ for\ Air\ Cooled\ Solution} = \frac{4U\ liquid\ cooled\ system}{10U\ air\ cooled\ system} = 40\%$$

$$WUE_{DLC\ vs\ AC} = \frac{Water\ for\ Liquid\ Cooled\ Solution}{Water\ for\ Air\ Cooled\ Solution} = 60\%$$

$$\$UE(MUE)_{DLC\ vs\ AC} = \frac{TCO\ for\ Liquid\ Cooled\ Solution}{TCO\ for\ Air\ Cooled\ Solution} = 80\%$$

| Metric | LC vs. AC Ratio | Savings |
|---|---|---|
| PUE (Power Usage Effectiveness) | **~60%** | **40%** |
| SUE (Space Usage Effectiveness) | **~40%** | **60%** |
| WUE (Water Usage Effectiveness) | **~60%** | **40%** |
| $UE (Money Usage Effectiveness) | **~80%** | **20%** |

**Note:** The above modelling is for a generic case; each data center design is unique and considers various factors such as climate, including seasonal temperatures and humidity. The above calculation illustrates the scale of savings across key metrics.

**PUE (Power Usage Effectiveness) Ratio:** The power consumption for a liquid-cooled data center divided by the power consumption for an air-cooled data center. Examining this definition in detail reveals the effectiveness of a liquid-cooled data center. Assumptions for this calculation are as follows: the traditional PUE for an air-cooled data center is assumed to be 1.5, while for a liquid-cooled data center, it is assumed to be 1.1. Therefore, the ratio of liquid-cooled data center power usage compared to that of the air-cooled data center would be 60%, resulting in savings of 40%.

**SUE (Space Utilization Effectiveness) Ratio**: Liquid-Cooled Rack Density / Air–Cooled Rack Density. This metric shows how much the computing density increases with liquid-cooled servers compared to air-cooled servers in a rack and thus a data center. Servers that use liquid cooling are designed to be more compact, thereby the quantity of racks and cables can be reduced significantly. Assumptions for this calculation:

- The height of an air-cooled server is 10U
- The height of a liquid-cooled server is 4U

Thus, from a GPU density perspective, a liquid-cooled rack is about 60% more dense than an air-cooled rack solution.

**WUE (Water Usage Effectiveness) Ratio**: is the amount of water saved when using liquid cooling in a data center compared to using air cooling. In an air-cooled data center, water cools the HVAC and the IT system. In a liquid-cooled system, water is used to cool the IT systems directly. More efficiency can be achieved by using warm water instead of pre-chilled water to cool systems. Our measurements suggest that about 40% less water is used when using liquid cooling by eliminating the use of chillers.

**$UE (Dollar or Money Usage Effectiveness) Ratio:** Compares the TCO of operating an air-cooled data center with a direct-to-chip liquid-cooled data center. A data center built with Supermicro's DLC-2 technology will use significantly less power and water (as illustrated above) while occupying a smaller footprint than a traditional air-cooled Data Center. These OpEx savings will typically lead to a 20% lower TCO over time.

## DLC-2 Data Center Improvements

A key component of the Supermicro DLC-2 offering is the inclusion of a new liquid-cooled optimized GPU server, the SYS-422GS-NBRT-LCC. This innovative GPU server is designed to reduce power consumption and decibel levels through the strategic placement of components & innovative design.  Key specs shown below:

- 4U Form Factor
- 2x Intel® Xeon® 6 Series processor with P-Cores (up to 500W per CPU)
- NVIDIA HGX with 8x B200 GPUs (up to 1,000W per GPU)
- Liquid cooling for GPUs, CPUs, DIMMS, PCIe Switches, VRMs, PWS (Power Supplies)
- Front I/O design for better thermal performance

*Figure 1 - SYS-422GS-NBRT-LCC*

Other Components of Supermicro DLC-2 Solution include:

- In-Rack CDU Enhancements – New in-rack CDUs can remove 250kW of heat generated per rack, a 150% increase over existing 100kW CDUs.
- Higher Inlet Temperature Servers – The temperature of the water supplied to the server can be increased. The new servers can use an inlet temperature of up to 45° C, which reduces the power needed to cool the liquid via a chiller.
- Hybrid cooling towers – New cooling towers bring together the features of standard and dry water towers into a single design. This is especially beneficial in locations with substantial seasonal temperature variation.
- Additional Cold Plates – New cold plates covering the GPUs, CPUs, PCIe switches, DIMM memory, voltage regulators, and power supply units (PSU). These will remove up to 98% of the heat generated by the system's components.
- Air Flow Reduction – With fewer components to cool by air, the fans within the system can run at lower speeds, reducing power consumption.
- Noise Generation – Lower fan speed results in lower decibel levels for each server, up to 25dB quieter.

| DLC-2 Key Metrics | |
|---|---|
| Cold Plates | CPU, GPU, PCIe Switch, DIMM, VRM, PSU |
| Max Inlet Water Temperature | Up to 45° C |
| Noise Level | As low as 50dB |
| DLC Server Heat Capture | Up to 98% |
| Server Power Savings (vs AC) | Up to 20% |
| Data Center Power Savings (vs AC) | Up to 40% |

## Conclusion

DLC-2 from Supermicro reduces data center's overall power consumption through the use of several innovative technologies, including a new optimized server design and hybrid cooling towers.

Liquid-cooled data centers offer numerous advantages, making them appealing choices for AI workloads and modern enterprise computing. Supermicro's Data Center Building Block Solutions (DCBBS) is an optimized end-to-end solution

comprising servers, racks, network, storage, water piping/tower, data center software management, and service, delivering the best time-to-market in the industry.

These systems enhance energy efficiency and reduce operational costs by significantly reducing power usage. The denser design of liquid-cooled rack-scale solutions also improves space utilization, allowing for more equipment within the same footprint. Additionally, the decrease in water usage compared to traditional air-cooling methods makes liquid cooling a more sustainable option while reducing OPEX and TCO. Overall, combining these benefits results in more cost-effective, environmentally friendly, and efficient data center operations regardless of location.

More importantly, Supermicro's DCBBS enables our clients to construct liquid-cooled data center infrastructure with the fastest time-to-online (TTO) deployment advantage, typically shortening a 12-month deployment process to just 3 months. This means that Supermicro provides an advantage to our customers by accelerating the return on their infrastructure investments sooner.

## For More Information

Supermicro Liquid Cooling Information: http://www/supermicro.com/liquidcooling

Supermicro.com Information: www.supermicro.com